

AD-A260 644



**DTIC**  
**ELECTE**  
**FEB 11 1993**  
**S C D**

University  
of Southern  
California



11

**DISTRIBUTION STATEMENT A**  
Approved for public release;  
Distribution Unlimited

INFORMATION  
SCIENCES  
INSTITUTE



310/822-1511

4676 Admiralty Way/Marina del Rey/California 90292-6695

Statement A per Mr. St. Johns, DARPA  
Code CSTO.  
Arlington, VA 22203

2/10/93 JK

Accession For	
NTIS	<input checked="" type="checkbox"/>
DTIC	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

**Gigabit Network Communications Research**  
**Quarterly R&D Status Report No. 8**

Period: 1 JUL 92- 30 SEP 92

DTIC QUALITY INSPECTED 3

*Sponsored by*  
*Defense Advanced Research Projects Agency (DoD)*  
*Computer Systems Technology Office*  
*Gigabit Network Communications Research*  
*Under Contract #DABT63-91-C-0001*  
*PR&C: HR0011-0218-0001*  
*AAP No. DAR1010*  
*Issued by Directorate of Contracting*  
*Fort Huachuca, AZ*

*The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.*

93-00760

93 1 13 003

# **GIGABIT NETWORK COMMUNICATIONS RESEARCH**

## **QUARTERLY R&D STATUS REPORT for DABT63-91-C-0001**

**(covering the period 7/1/92-9/30/92)**

### **1. ATOMIC**

The ATOMIC project utilizes inexpensive fine-grain multicomputer components to create a multi-gigabit per second local area network. This is a unique approach that leverages recent DARPA-supported research at Caltech by Professor Charles Seitz and his students. That research has already produced dense and extremely fast Mosaic single-chip multicomputer nodes. ATOMIC is a unique example of network development that grows upwards from multicomputer machine networking into the domain of a LAN.

The principal concerns for this quarter were the continued integration of the new SPARCstation ATOMIC SBus card, modification of the Address Consultant to recognize SBus hosts, and testing of the 3x3 ATOMIC mesh.

#### **1.A. SOFTWARE**

To create a new LAN requires a large amount of software to be developed. We continue to make extensions to the BSD UNIX kernel to enhance the operation of the LAN.

ATOMIC SBus device driver and kernel development to attach the SBus cards to SUN SPARC platforms has been finished. Packets began moving across the network from SPARCs in September.

ATOMIC device driver and Mosaic programs have been finished that enable Sun SPARCstations to be connected to the ATOMIC LAN via the new SBus host interface cards that were developed at Caltech. In addition, we converted our various performance tools to work with the latest version of the ATOMIC kernel.

Some Address Consultant and kernel IOCTL call development remains to be finished. This will be completed in October 1992.

#### **1.B. HARDWARE**

##### *1.B.1. ATOMIC Switch and Mosaic Chips*

The ATOMIC Project has demonstrated a multi-gigabit per second crossbar or concentrator switch. The first steps were taken last quarter. The first prototype ATOMIC mesh switch, a 3x3 mesh, arrived in the last week of June. We reported last quarter that packets had been routed through the mesh.

This quarter, the prototype 3x3 mesh's performance was characterized. It has run reliably for the past three months. In the diagram produced by the Address Consultant

display driver (see below) you can see six hosts, including two SPARCstation hosts, that are connected to the 3x3 mesh.


Sending six streams between hosts through the mesh has demonstrated that the switch easily handles loads well in excess of 1.3 Gb/s with no signs of loading. For these tests only 1/2 the switch input and output ports were used and each port was driven at less than 50% capacity. This supports our conclusion, drawn from Caltech simulation studies, that a 3x3 ATOMIC mesh switch will route at least 4 Gb/s of traffic. The single board 8x8 switch should handle a 10 Gb/s load.

**Measured performance:**

1.3 Gb/s (6 streams at 210 Mb/s)  
4 Mpkt/s (8 streams at 0.5 Mpkt/s)

There was some frustration integrating the switch into the system due to incompatibilities in the Host Interface channel drivers and the mesh channel drivers. In our zeal to move forward we have taken some engineering shortcuts regarding the electrical aspects of the channel signals. We have a stable configuration, but we should be buffering the mesh channel signals. This is an issue that Caltech is solving for us, so there is little point in our duplicating that engineering.

We have completed the non-consistent mapping and routing version of the AC. We can now map and route over any legal ATOMIC/Mosaic interconnection topology. Work on optimizing the AC also occurred. The AC mapping algorithm to reduce the number of required messages by an order of magnitude.



### *1.B.2. Active Cable Progress*

Work began on the design of a prototype fiber-optic 30m-kilometer ATOMIC channel in late September and will be presented in the next quarterly report. AT&T 41mm series parts to drive 25-pair phone cable are being examined for use as a less expensive alternative to fiber-optics when distances shorter than 30 meters are acceptable. They are being characterized in October. The use of HP HDMP 1002/1004 parts are being considered for the fiber-optic version of the ATOMIC channel.

## **2. Personal Conferencing**

The implementation of the Connection Control Protocol (CCP) in the Multimedia Conference Control program (MMCC) was tested over DARTnet this quarter. Both the specification and the implementation were further refined based on the test results.

To better understand how CCP and the connection management architecture of which it is part may be used in applications, we have held several teleconference discussions with other researchers at BBN and SRI who are exploring the use of CCP for multimedia synchronization, collaborative design and Simnet host control. An alpha release of the CCP sources was made available to these researchers. We are also comparing this architecture with that of the Bellcore Touring Machine system being installed at DARPA under ISI's ACT project.

We are participating in IETF working groups to develop broader solutions for telecollaboration across the Internet. At the July IETF in Boston, Eve Schooler gave a presentation about CCP in the Remote Teleconferencing BOF. In the Audio/Video Transport working group, Steve Casner presented and led discussion on the second iteration of a strawman for a real-time transport protocol.

To test theory with practice, we helped coordinate live transmission of packet audio and video from these working group meetings and the IETF plenary sessions across the Internet to remote participants on 170 different hosts in 10 countries. This is a critical test to see how well IP multicast protocols will scale up for wider deployment of teleconferencing in the Internet. Subsequent to the IETF meeting, we have played a major role in the coordination and construction of the MBONE, a virtual IP multicast backbone network to support future IETF audio/videocasts and other experiments with an even larger number of participants and countries.

One means to enable packet video experiments on a larger scale is to implement software decoding of compressed video to allow a large number of receivers without additional hardware cost. Working with students at USC, we integrated routines written last quarter to decode compressed video from the Concept/Bolter video codecs used in DARTnet experiments with packet reception and X-window display software. This would allow the DARTnet sites without codec hardware to join in experiments and teleconferences at least as receive-only sites. It also facilitates workstation-based personal conferencing by eliminating the need for a separate monitor on which to display video.

### **3. Integrated Services Protocols**

In July 1992, we delivered an address to the plenary of the Internet Engineering Task Force meeting, entitled "DARPA Research Network Testbed - Progress Report". This talk emphasized the work in DARTnet towards development of realtime service, including Integrated Service protocols.

We began the formulation of a new research plan for DARTnet, including (1) completion of the 'DARTnet teleconferencing technology' and (2) a quasi-production mode of operation.

Work in this area is continuing.

### **4. IP/SQ Congestion Control**

Data has been gathered from DARTnet experiments. The IP/SQ algorithm operates approximately as expected. Using the TG program to generate UDP test streams, single source testing indicates that IP/SQ can reduce congestion loss from 70% to approximately 16%.

The use of deep queues in the HSI output drivers in conjunction with the use of random-drop by IP/SQ causes variance in the distribution of round-trip times seen by a source in returning Source Quench messages. This variance becomes more of an issue for short round-trip times. Round-trip times as short as 29 milliseconds have been observed during testing.

To adjust for latency, a modification in the algorithm has been made. The time-stamp that the source inserts in the header of outgoing IP packets is extended to include a queue-residency-time field. That field is time-stamped on arrival at a gateway or destination. When that site sends a Source Quench it inserts in that field the elapsed time spent in the queue.

That elapsed time is then used to decrement the apparent round-trip time derived from the source timestamp when the source receives the Source Quench message. This improves the accuracy of the round-trip time to the point of congestion, allowing the algorithm to become more aggressive but to do so still with a safely chosen feedback interval.

The conversion of DARTnet from PSI to Sprint interrupted testing for several weeks. Debugging and testing with the improved round-trip estimator will begin shortly.

### **5. Tunnel**

The Tunnel Project is in the process of acquiring equipment that is needed in order to start work on development. A SPARCstation MP, which is on loan from DARPA, was installed on September 16th. A second Ethernet interface board has been ordered, and is

expected to arrive on October 9th. The Tunnel Project is also borrowing an IBM PC, which belongs to ISI's Division 7. The PC will be connected to the SPARCstation via the second Ethernet interface and will be used to simulate a restricted network.

## **6. Automated Cluster Teleconferencing (ACT)**

The Touring Machine (TM) port to the NeXT using Co-Xist has been completed. The port uses MIT's software to control the AutoPatch switch. Single-domain, multi-user conferences are supported. The TM was modified to support bridging without Bellcore's proprietary interface. It was decided not to make a native port to NeXTStep Windows, as that would reduce the generality of the ACT software for future uses.

Runtime modules of Co-Xist were ordered for DARPA. Analog audio and video alternatives were evaluated, but were only marginally beneficial and would have delayed installation. We will use the originally specified analog audio and video equipment. An alternate camera is under evaluation while availability is determined. The ConferTech audio bridge was substituted by a less expensive TEAC studio mixer. The audio bridging of the TEAC was configured to correspond to the video mixing of the Panasonic unit. The switching and bridging hardware assembly has been configured and cabled.

The next quarter will include installation of the current system, development and installation of the multi-domain extension, and determination of the second half goals of the ACT project. Installation of the current system is projected for early December, pending delivery of video equipment and office cabling.

Equipment for DARTnet teleconferencing (for ACT project management use) will be installed at DARPA shortly (approximately 8 OCT 92).

## **7. DARTnet Network Operations Center (DARTNOC)**

### **7.A. Operations**

ISI continues to schedule test times for experimenters.

ISI and others supported the audio/video conferencing for the Boston July 92 IETF.

ISI coordinated the installation of the DARTnet new vendor circuits. ISI has also developed a cut-over plan to minimize the disruption caused by moving to the new vendor supplied network. Installation is scheduled for next month.

### **7.B. Hardware/Software**

ISI prepared disks with 4.1.1 software, cleaned up configurations, and repartitioned disks for installation in routers at NASA/Ames, LA-POP (soon to be Anaheim), DC-POP, ISI, and DARPA. These sites will have either two 100 MB disks or one 200

MB disks with two 100 MB partitions. This will permit us to install two independent operating systems on these disks. One might be Sun OS based and one BSD 4.4 based. We will also be able to install new software for an OS upgrade over the network without fear of overwriting the files needed to complete the upgrade. That is, we will be able to completely install a new OS version and test it before discarding the old version. Other sites will be upgraded as soon as we can get more disks.

## **8. Infrastructure**

### **8.A. USER SERVICES**

As the USAC Chair, Joyce Reynolds participated in IESG Teleconferences from July – September 1992 and attended the IETF meeting in Boston, Mass., July 13–17, 1992.

Eight working groups in the User Services Area of the IETF met in Boston, MA. One BOF (Birds of a Feather) was held in the User Services area regarding a working group formation on WHOIS and User Lookup Services.

During this period, four new Working Groups have been formed in the User Services Area of the IETF:

- Integration of Internet Information Resources (IIIR)
- Networked Information Retrieval (NIR)
- Universal Resource Identifiers (URI)
- Whois and Network Information Lookup Service (WNILS)

There are currently 13 active Working Groups in the User Services Area of the IETF.

Two papers were published as FYI RFCs:

- FYI 16        "Connecting to the Internet: What Connecting Institutions Should Anticipate", (Also RFC 1359), August 1992.
- FYI 15        "Privacy and Accuracy Issues in Network Information Center Databases", (Also RFC 1355), August 1992.

### **8.B. INTERNET MONTHLY REPORT**

The Internet Monthly Report (IMR) is the status report on the operation of the Internet and the research and development activities of the Internet community. It features reports from the IAB, the Internet Research Task Force and its research groups, and the Internet Engineering Task Force and its working groups in addition to the reports from approximately 30 regional networks and individual sites. A typical monthly report is approximately 40 pages.

During this reporting period, three Internet Monthly Reports for June 1992, July 1992, and August 1992 were assembled, edited, and distributed directly (via electronic mail)



to over 375 mailboxes, some of which are exploder mailboxes where the report is sent to a list of people. In particular, the mailbox "IETF@isi.edu", which is one of the mailboxes on the IMR list, goes out to an additional 935 mailboxes, many of which are further exploders.

## **8.C. HIGH PERFORMANCE NETWORK RESEARCH REPORT**

The High Performance Network Research Report (HPNRR) discusses research and development activities in the Gigabits program and the advanced networking research community. A typical report is about 25 pages. During this reporting period, three reports for June 1992, July 1992, and August 1992 were assembled, edited, and distributed directly (via electronic mail) to over 130 people.

## **8.D. REQUEST FOR COMMENTS**

ISI serves as the technical editor and publisher of the Internet document series called "Requests for Comments" (RFC). 3 RFCs were published this quarter:

- RFC 1340: Reynolds, J. K., and J. Postel, *Assigned Numbers*, ISI, July 1992.
- RFC 1348: Manning B., *DNS NSAP RRs*, Rice University, July 1992.
- RFC 1349: Almquist, P., *Type of Service in the Internet Protocol Suite*, Consultant, July 1992.
- RFC 1350: Sollins, K., *THE TFTP PROTOCOL (REVISION 2)*, [STD 33], MIT, July 1992.
- RFC 1351: Davin, J. (MIT), J. Galvin (TIS), and K. McCloghrie (Hughes LAN), *SNMP Administrative Model*, July 1992.
- RFC 1352: Galvin, J. (TIS), K. McCloghrie (MIT), and J. Davin (MIT), *SNMP Security Protocols*, July 1992.
- RFC 1353: McCloghrie, K. (Hughes LAN), J. Davin (MIT), and J. Gavin (TIS), *Definitions of Managed Objects for Administration of SNMP Parties*, July 1992.
- RFC 1354: Baker, F., *IP Forwarding Table MIB*, ACC, July 1992.
- RFC 1355: Curran, J. (NNSC), and A. Marine (SRI), *Privacy and Accuracy Issues in Network Information Center Databases*, August 1992.
- RFC 1356: Malis, A. (BBN), D. Robinson (Computervision Systems Integration), and R. Ullman (Process Software Corp.) *Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode*, August 1992.

- RFC 1358: Chapin, L., Chair (IAB), *Charter of the Internet Architecture Board (IAB)*, August 1992.
- RFC 1359: ACM SIGUCCS Networking Taskforce, *Connecting to the Internet--What Connecting Institutions Should Anticipate*, August 1992.
- RFC 1360: Postel, J. (IAB), *IAB Official Protocol Standards*, September 1992.
- RFC 1361: Mills, D., *Simple Network Time Protocol (SNTP)*, August 1992.
- RFC 1362: Allen, M., *Novell IPX Over Various WAN Media (IPXWAN)*, Novell, Inc., September 1992.
- RFC 1363: Partridge, C., *A Proposed Flow Specification*, BBN, September, 1992.
- RFC 1364: Varadhan, K., *BGP OSPF Interaction*, OARNET, September 1992.
- RFC 1365: Siyan, K., *An IP Address Extension Proposal*, Siyan Consulting Services, September 1992.

#### **8.E. VISITORS**

Jon Postel and Peter Will hosted a Gigabit meeting with guests Ira Richer and Mike St. Johns, August 5, 1992.

#### **8.F. TRAVEL**

Bob Braden, Jon Postel, Steve Casner, Eve Schooler, and Joyce Reynolds, attended the IETF meeting in Boston MA, Jul 14-17, 1992.

Robert Felderman gave a talk at the Technology Transfer Institute entitled, "Gigabit Networks Seminar" in Washington DC on July 28.

Peter Will, Danny Cohen, and Jon Postel attended the DARPA PI meeting in New Orleans September 1-2, 1992.

#### **8.G. SEMINARS**

BOB FELDERMAN discussed "GIGABYTES to GIGA-SITES" and recent work on the ATOMIC project. Bob covered mapping and routing for an arbitrary topology and cycle detection for deadlock prevention. This was followed by an exploration of the future of the project centering mostly on the Address Consultant and it's ability to provide QOS guarantees.

## 9. Publications, papers, and presentations

Reynolds, J. K., and J. Postel, *Assigned Numbers*, RFC 1340, ISI, July 1992.

Cohen, D., G. Finn, R. Felderman, A. DeSchon, *ATOMIC: Low-Cost, Very High-Speed LAN*, ISI/RR-92-291, September 1992.

Postel, J. (IAB), *IAB Official Protocol Standards*, RFC 1360, September 1992.

We have submitted a paper on the novel approach taken in creating the ATOMIC LAN to the Symposium on Integrated Systems to be held at the University of Washington. It is entitled: "A VLSI Approach to Local Area Communication".

A paper on "Physics Analogs in Communication Models" was submitted and accepted at the Physics of Computation Workshop (10/92).

A paper entitled "Communication Parallelism" was submitted to IEEE InfoCom '93; notification is due mid-October.

A letter to IEEE Communications Magazine, regarding an article on gigabit network protocol issues, has been revised and is awaiting submission via an external co-author.

An article on "Replication and Reduction in Multistage Interconnection Networks" for IEEE Computer has been tabled, pending available resources for revision.